REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-09-0037

1. REPORT DATE *(DD-MM-YYYY)* 10-02-2009	**2. REPORT TYPE** Final Report	**3. DATES COVERED** *(From - To)* July 2007 - November 2008
4. TITLE AND SUBTITLE Performance Optimization of Complex Systems		**5a. CONTRACT NUMBER**
		5b. GRANT NUMBER FA9550-07-1-0213
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Cassandras, Christos G.		**5d. PROJECT NUMBER**
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Boston University, Boston, MA 02215		**8. PERFORMING ORGANIZATION REPORT NUMBER** N/A
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NL 875 North Randolph St Arlington, VA 22203 Dr. Robert Hearn		**10. SPONSOR/MONITOR'S ACRONYM(S)** AFOSR
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution A: Approved for Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The research results in this report are based on an effort to develop systematic techniques for performance optimization of complex systems with special emphasis on real-time methods. Recognizing the increasing importance of stochastic networks in both the civilian and military domains, explicit algorithms are sought that are scalable, distributed, asynchronous, and computationally compatible with the limited processing capabilities at individual nodes of many such networks. The main outcomes of the project are: (a) An asynchronous event-driven distributed optimization framework allowing autonomous agents to cooperate toward a common goal with minimal communication among them, thus saving energy without any loss in performance. In particular, communication is limited to instants when a state estimation error function at some agent exceeds a threshold. (b) An optimization framework for systems with time-critical tasks in which hard real-time constraints are guaranteed to be satisfied. At the single node level, an efficient solution procedure termed the Critical Task Decomposition Algorithm (CTDA) was developed. At the multi-node level, a Virtual Deadline Algorithm (VDA) was developed. Both algorithms are scalable in the number of tasks executed. (c) Extensions to perturbation analysis methods for gradient estimation and optimization of Stochastic Fluid Models (SFM) as abstractions of complex stochastic systems.

15. SUBJECT TERMS

Distributed Optimization Algorithms, Asynchronous Optimization Algorithms, Autonomous Agents, Cooperative Systems, Time-Critical Systems, Perturbation Analysis, Stochastic Fluid Models

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** Christos G. Cassandras
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	UU	9	**19b. TELEPHONE NUMBER** *(include area code)* 617-353-7154

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

FINAL REPORT

AFOSR GRANT: FA9550-07-1-0213

TITLE: **Performance Optimization of Complex Systems**

REPORTING PERIOD: July 2007 – November 2008

Principal Investigator:

Christos G. Cassandras
Division of Systems Engineering
Boston University
Boston, MA 02215
Tel: (617) 353-7154
FAX: (617) 353-4830
e-mail: cgc@bu.edu

February 2009

1. OBJECTIVES

The research pursued under grant FA9550-07-1-0213 has been aimed at developing systematic techniques for optimizing the performance of complex systems and exploring some new innovative directions for real-time optimization methods. Recognizing the increasing importance of stochastic networks in both the civilian and military domains, special emphasis has been placed on techniques and explicit algorithms that are scalable, implementable in distributed and asynchronous fashion, and computationally compatible with the limited processing capabilities at the individual node level of many such networks.

Three major research objectives have been pursued:

1. Developing *asynchronous*, as well as *distributed*, optimization algorithms for systems consisting of multiple autonomously operating "agents" (UAVs, nodes in a sensor network, etc) that must cooperate so as to optimize a common objective. An "asynchronous" approach allows each agent to act and communicate with other agents on an event-driven basis, not constrained by a synchronizing clock. This has significant practical benefits, as explained in the next section.

2. Developing optimization algorithms for systems required to perform time-critical tasks. These systems often consist of components with limited computational resources, e.g., small, wireless, battery-dependent devices. Therefore, the challenge is to develop fast, scalable solution methods for such optimization problems, as opposed to conventional "brute force" approaches suitable only for off-line applications.

3. Developing Stochastic Fluid Models (SFM) as hybrid system abstractions of complex stochastic systems and explore the use of perturbation analysis and gradient-based optimization methods for such models.

2. ACCOMPLISHMENTS AND NEW FINDINGS

1. Asynchronous Distributed Cooperative Optimization with Minimal Communication.

Distributed optimization is necessary in systems consisting of multiple autonomously operating components (UAVs, nodes in a communication or sensor sensor network, etc) that must cooperate so as to optimize a common objective. This requires communication among the networked components so as to exchange local state information. The cost of such communication is significant, especially when the system components are wireless devices with limited energy. Therefore, we have addressed the basic question: "What is the minimal amount of communication required in a cooperative system for an optimization problem to still be solvable with no loss of its fundamental optimality properties?" To address this question, we have sought conditions under which communication of state information among nodes can be minimized while still ensuring that the optimization process converges. This necessitates developing *asynchronous* event-driven optimization algorithms which allow each system component to operate autonomously, not constrained by clock synchronization, so that it acts or communicates with other components only whenever it detects an event that warrants such activity.

We have developed an asynchronous distributed optimization framework that limits communication exclusively to instants when a state estimation error function at some

component exceeds a threshold [10]. Thus, a component communicates its local state to others on an event-driven basis and only "as a last resort". We have formally proved that, under certain technical conditions, such convergence is guaranteed, as reported in [16]. The implication is that cooperation toward a common objective may be achieved with minimal communication (possibly none at all) among system components. We have applied this approach to a sensor network coverage optimization problem where the objective is to maximize the probability of detecting events occurring in a given region [16]-[17]. We have shown that optimality is achieved with limited asynchronous communication among nodes at a fraction of the cost required under periodic state information exchange.

We have also developed an interactive simulation environment, available on the web at *http://codescolor.bu.edu/real%20time%20coverage/coverage_demo.html*, through which we have performed experiments and tests related to the asynchronous distributed optimization framework described above and compared it to standard synchronous approaches. In addition, we have created a laboratory setting with small wireless mobile robots (Khepera III) acting as "agents" in a "real-world" cooperative setting (see Figure 1 and the web site *http://codescolor.bu.edu/multimedia.html*).

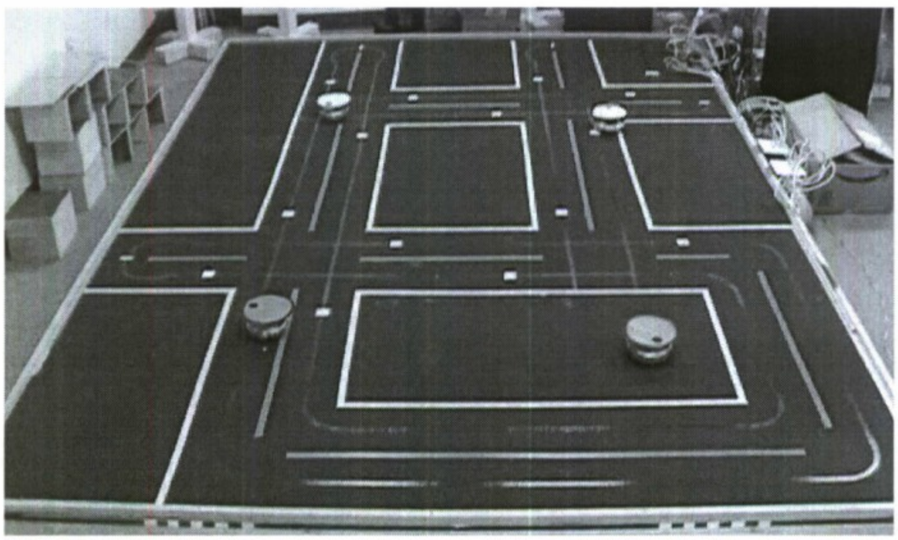

Figure 1: Laboratory setup at Boston University for cooperative optimization

- *Relationship to original goals*: This work is within the scope of Objective 1 (see Section 1).

- *Relevance to AF mission and potential applications:* As systems with small, inexpensive, wireless devices become increasingly more prevalent, dramatic improvements are needed in optimization schemes embedded in and executed by such devices. This is achievable through distributed, highly efficient algorithms, but also, perhaps more importantly, asynchronous in nature. Asynchronous (in particular, event-driven) operation is crucial for military systems in environments that are resource-limited and possibly adverse for the following reasons: (*i)* Saving energy and hence prolonging system lifetime, especially in a wireless setting, (*ii*) Reducing security risks (jamming, data corruption, etc) since such operation minimizes the need for communication, (*iii*) Overcoming the requirement for clock synchronization, which is generally infeasible, (*iv*) Overcoming bandwidth limitations.

2. Scalable Optimization of Systems with Real-Time Constraints

Systems required to perform time-critical tasks include communication networks, sensor networks, and UAV teams executing time-sensitive missions. This gives rise to difficult optimization problems with nonlinear objective functions and nondifferentiable inequality constraints that model the time-critical task execution dynamics. In addition, these systems often consist of components with limited computational resources, e.g., small, wireless, battery-dependent devices called upon to carry out optimization-based decision making processes. Therefore, the challenge is to develop fast, scalable solution methods for such optimization problems, as opposed to conventional "brute force" approaches suitable only for off-line applications.

At the single node (or single stage) level, we have formulated the basic optimization problem with real-time constraints and shown that its solution is characterized by several attractive structural properties leading to an efficient solution procedure termed the *Critical Task Decomposition Algorithm* (CTDA) [4]. At the multi-node level, corresponding to a series of tasks with end-to-end real-time deadlines, we have shown that the optimal solution is equivalent to one with "virtual" deadlines imposed at each node. Based on this property, we have developed a *Virtual Deadline Algorithm* (VDA) that takes advantage of the efficiency of the CTDA, as described in [2]. We have also proved that the VDA converges to the global optimum for the original problem and that, just like the CTDA, it is scalable in the number of tasks executed. We have also been able to further extend the VDA to a multi-layer network architecture [7].

The above results are based on the assumption that there exists an optimal solution such that tasks to be performed can all meet assigned hard real-time constraints. This leads to the question: "How can we deal with problems where it is impossible to perform all tasks within these constraints?" We have approached this problem in two ways. The first is by exercising a task "admission" policy: We seek to maximize the number of tasks that can be processed with guaranteed real-time constraints (and reject the rest). This seemingly combinatorially hard problem can in fact be solved through a fast and scalable algorithm under certain conditions that are satisfied in most practical cases of interest; details are reported in [11]. The second way is by replacing the hard deadline requirements by "weekly hard" ones where *m* out of any *k* consecutive tasks must meet their respective deadlines. We have shown in [3] that the CTDA may be extended to accommodate this case while maintaining its scalability properties.

The problem of optimization with real-time constraints becomes much more challenging in a stochastic setting where "tasks" to be performed by a node are not known in advance. We have, therefore, initiated a study of such a problem where the node must react to every new task and possibly use statistical information about its environment and the traffic it handles. We have carried out a "worst-case" analysis, which may be conservative. With this motivation, we have subsequently explored a drastically different idea that exploits the type of tasks and performance criteria we are interested in. This new approach provides a "best solution in probability" efficiently obtained by estimating the probability distribution of sample-path-optimal solutions. Early results [12] show substantial performance improvements over worst-case analysis.

• *Relationship to original goals*: This work is within the scope of Objective 2 (see Section 1).

• *Relevance to AF mission and potential applications:* Systems with time-critical tasks arise in both military and civilian settings. Examples include military missions whose goals must be attained within rigid deadlines, and health-related applications where patient data must be retrieved and processed by certain deadlines. Satisfying hard real-time constraints poses a challenge as it significantly complicates otherwise manageable optimization problems and may often result in posing infeasible problems, hence unnecessarily expending resources to seek nonexistent solutions.

3. Perturbation Analysis and Optimization of Stochastic Flow Models (SFM)

Recognizing the importance of stochastic networks for communication, transportation, and UAV missions, we have continued our past work studying a class of hybrid systems termed Stochastic Fluid Models (SFM) that represent complex stochastic networks, replacing detailed discrete-event models which become practically infeasible. Under past AFOSR grants, we combined perturbation analysis and gradient-based optimization techniques to develop network management algorithms that are easily implementable. In this project, we were able to extend this work to systems whose performance metrics include constraints on task delays (along the lines of the real-time constraint problems discussed earlier, but in a stochastic setting), rather than the traditional average delay metric. We were able to develop new gradient estimators and establish their unbiasedness in [5].

• *Relationship to original goals*: This work is within the scope of Objectives 2 and 3 (see Section 1).

• *Relevance to AF mission and potential applications*: Stochastic networks manifest themselves in communication, transportation, and UAV applications. Their complexity precludes the use of traditional discrete event models and calls for abstractions of sufficient accuracy to enable the development of performance optimization algorithms.

3. PUBLICATIONS RESULTED FROM FA9550-07-1-0213

• Papers Published:

[1] Cassandras, C.G., and Li, W., "Optimization Problems in the Deployment of Sensor Networks", in *Modeling and Control of Complex Systems* (P. Ioannou and A. Pitsillides, Ed's), pp. 179-202, 2007.

[2] Mao, J., and Cassandras, C.G., "Optimal Control of Multi-Stage Discrete Event Systems with Real-Time Constraints", *IEEE Trans. on Automatic Control*, AC-54, 1, pp. 108-123, 2009.

[3] Zhuang, S., and Cassandras, C.G., "Optimal Control of Discrete Event Systems with Weakly Hard Real-Time Constraints", *J. of Discrete Event Dynamic Systems*, Vol. 19, 1, pp. 67-89, 2009.

[4] Mao, J., and Cassandras, C.G., "Scalable Optimization Algorithms for Discrcte Event Systems with Real-Time Constraints: An Overview of Recent Developments", *Proc. of 2008 Intl. Workshop on Discrete Event Systems*, pp. 150-155, May 2008.

[5] Panayiotou, C.G., and Cassandras, C.G., "IPA for Delay Threshold Violation Using Stochastic Fluid Models", *Proc. of 2008 Intl. Workshop on Discrete Event Systems*, pp. 162-167, May 2008.

[6] Zhong, M., and Cassandras, C.G., "Distributed Coverage Control in Sensor Network Environments with Polygonal Obstacles", *Proc. of 17th IFAC World Congress*, pp. 4162-4167, July 2008.

[7] Mao, J., and Cassandras, C.G., "Optimal Control of Multi-layer Discrete Event Systems with Real-Time Constraint Guarantees", *Proc. of 17th IFAC World Congress*, pp. 4120-4125, July 2008.

[8] Ning, X., and Cassandras, C.G., "Optimal Dynamic Sleep Time Control in Wireless Sensor Networks", *Proc. of 47th IEEE Conf. Decision and Control*, pp. 2332-2337, Dec. 2008.

[9] Mao, J., and Cassandras, C.G., "On-line Optimal Control of a Class of Discrete Event Systems with Real-Time Constraints", *Proc. of 47th IEEE Conf. Decision and Control*, pp. 1471-1476, Dec. 2008.

[10] Zhong, M., and Cassandras, C.G., "Asynchronous Distributed Optimization with Minimal Communication", *Proc. of 47th IEEE Conf. Decision and Control*, pp. 363-368, Dec. 2008.

• **Accepted, but not yet published:**

[11] Mao, J., and Cassandras, C.G., "Optimal Admission Control of Discrete Event Systems with Real-Time Constraints", to appear in *J. of Discrete Event Dynamic Systems*, 2009.

[12] Mao, J., and Cassandras, C.G., "On-line Optimal Control of a Class of Discrete Event Systems with Real-Time Constraints", to appear in *J. of Discrete Event Dynamic Systems*, 2009.

• **Submitted, but not yet accepted:**

[13] Wesselowski, K., and Cassandras, C.G., "A Cooperative Receding Horizon Approach to the Elevator Dispatching Problem", subm. to *IEEE Trans. on Automation Science and Engineering*, 2007.

[14] Ning, X., and Cassandras, C.G., "Message Batching in Wireless Sensor Networks - A Perturbation Analysis Approach", subm. to *J. of Discrete Event Dynamic Systems*, 2008.

[15] Ning, X., and Cassandras, C.G., "Dynamic Sleep Time Control in Wireless Sensor Networks", subm. to *ACM Trans. on Sensor Networks*, 2008.

[16] Zhong, M., and Cassandras, C.G., "Asynchronous Distributed Optimization with Minimal Communication", subm. to *IEEE Trans. on Automatic Control*, 2009.

[17] Zhong, M., and Cassandras, C.G., "Asynchronous Distributed Algorithms for Optimal Coverage Control with Sensor Networks", subm. to *17th IEEE Mediterranean Conference on Control and Automation*, 2009.

4. PERSONNEL SUPPORTED

• Principal Investigator:

Christos G. Cassandras, Professor, Boston University

• PostDoc:

Jianfeng Mao

• Graduate Students:

- Jianfeng Mao (PhD obtained, 2008)
- Chen Yao
- Minyi Zhong
- Ali Kebarighotbi

The PhD dissertation completed by Jianfeng Mao is entitled "Dynamic Energy Management in Resource Limited Systems with Real-Time Constraints". It considered problems motivated by the increasing need for energy management in resource limited systems with real-time constraints. Two techniques were utilized to minimize energy consumption: Dynamic Voltage Scaling and Dynamic Transmission Control, which can both greatly conserve energy by finely tuning the voltage level in a processor and the transmission rate in a transceiver unit. Since both techniques deal with a similar tradeoff between energy and latency, they were considered in the same context, treated as discrete event systems with hard real-time constraints, where tasks are non-preemptive and aperiodic. They were studied as dynamic optimization problems whose objective is to minimize energy consumption by dynamically tuning the processing rate (equivalent to the voltage level or the transmission rate) subject to real-time execution constraints. These problems were approached along three dimensions: topology, cost function and control framework. Starting with the single-stage problem, a highly efficient algorithm was developed, termed *Critical Task Decomposition Algorithm* (CTDA), which exploits the structural properties of an optimal state trajectory. This can solve the off-line problem without relying on any nonlinear programming solver. The related admission control

problem was also studied and a *Maximal Shift Task Algorithm* (MSTA) was obtained to guarantee feasibility for both off-line and on-line single-stage problems. Subsequently, the on-line version of the problem was tackled by using two methods: a worst-case analysis and a probabilistic comparison algorithm. The former can bypass the complexity of random effects and has low computational complexity. The latter can improve performance by using probability distribution information based on a novel idea providing an alternative to standard stochastic programming, in which the efficiency of the off-line CTDA solution algorithm can be fully utilized. For a two-stage system with homogeneous cost functions it was determined that the optimal structural properties exploited in single-stage systems are no longer applicable. After discovering new optimal structural properties, a novel efficient algorithm was derived, termed *Virtual Deadline Algorithm* (VDA), which can obtain the global optimal solution by solving one-dimensional bounded convex optimization problems in spite of the high dimensionality of the original problem. Finally, a multi-stage system and multi-layer system was considered with general cost functions. In this case, only some of the previous properties of the optimal state trajectory still apply, but they need to be appropriately redefined. Based on that, an efficient algorithm, still termed *Virtual Deadline Algorithm* (VDA), was developed, which can obtain the global optimal solution by solving low-dimensional bounded convex optimization problems independent of the dimensionality of the original problem.

5. INTERACTIONS/TRANSITIONS DURING REPORTING PERIOD

Participation/Presentations at Meetings, Conferences, Seminars

C.G. Cassandras gave invited talks/ plenary addresses/courses at the following meetings/organizations:

- University of Cyprus, Nicosia, Cyprus, October 2007 (Invited Seminar)
- BBN Technologies, Cambridge, MA, October 2007 (Invited Seminar)
- NSF Worshop on Systems and Control for High School Students, New Orleans, LA, December 2007 (Invited Talk)
- *47th IEEE Conf. on Decision and Control*, New Orleans, LA, December 2007 (Invited Talk)
- NSF EFRI Program meeting, Arlington, VA, December 2007
- *8th International Conference on Cooperative Control and Optimization*, January 2008, Gainesville, FL, (Plenary Address)
- AFOSR Grantee Meeting, Arlington, VA, April 2008
- *3rd International Conference on Innovative Computing, Information and Control,* Dalian, China, June 2008 (Plenary Address)
- *17th IFAC World Congress*, Seoul, Korea, July 2008 (Invited Talk)
- Workshop on Systems and Control for High School Students, Seoul, Korea, July 2008 (Invited Talk)
- *8th International Conference on Control*, Manchester, UK, September, 2008 (Plenary Address)
- *2008 INFORMS Annual Meeting*, Washington, DC, October 2008 (2 Invited Talks)

Consultative and Advisory Functions

• Mathworks, Inc.: C.G. Cassandras is developing discrete-event and hybrid simulation software, Contact person: Michael Clune

Transitions

• *Incorporating Discrete Event and Hybrid System capability in MATLAB/Simulink software suite, The MathWorks*

C.G. Cassandras has been working since 2004 with The MathWorks, Inc., producer of the MATLAB/Simulink software suite that dominates the scientific computation market. As of November 2005, Mathworks has released a new product (SimEvents) incorporating a Discrete Event System capability into their products. SimEvents includes capabilities of simulating a variety of commercial and military systems. Contact person: Michael Clune.

6. NEW DISCOVERIES, INVENTIONS, OR PATENT DISCLOSURES

None over the reporting period.

7. HONORS/AWARDS

C.G. Cassandras (Lifetime):

Lilly Fellow (1991), Fellow of IEEE (1996), Fellow of IFAC (2008),

IFAC Harold Chestnut Prize (1999),

Distinguished Member Award, IEEE Control Systems Society (2006),

Editor-in-Chief of *IEEE Transactions on Automatic Control* (1998-present).

Honors/Awards received during grant period:

C.G. Cassandras:

- Elected Fellow of the International Federation of Automatic Control (IFAC)